

Package ‘rdlocrand’

May 14, 2026

Type Package

Title Local Randomization Methods for RD Designs

Version 2.0

URL <https://rdpackages.github.io/>,
<https://github.com/rdpackages/rdlocrand>

BugReports <https://github.com/rdpackages/rdlocrand/issues>

Description Provides tools to perform randomization inference for RD designs under local randomization: `rdrandinf()` to perform hypothesis testing using randomization inference, `rdwinselect()` to select a window around the cutoff in which randomization is likely to hold, `rdsensitivity()` to assess the sensitivity of the results to different window lengths and null hypotheses and `rdrbounds()` to construct Rosenbaum bounds for sensitivity to unobserved confounders. See Cattaneo, Titiunik and Vazquez-Bare (2016) <https://rdpackages.github.io/references/Cattaneo-Titiunik-VazquezBare_2016_Stata.pdf> for further methodological details, and references.

Imports AER, sandwich

Depends R (>= 3.1)

License GPL-3

Encoding UTF-8

NeedsCompilation no

Author Matias D. Cattaneo [aut, cre],
Rocio Titiunik [aut],
Gonzalo Vazquez-Bare [aut]

Maintainer Matias D. Cattaneo <matias.d.cattaneo@gmail.com>

Config/roxygen2/version 8.0.0

Repository CRAN

Date/Publication 2026-05-14 15:30:02 UTC

Contents

rdlocrand-package	2
rdrandinf	3
rdrbounds	7
rdsensitivity	9
rdwinselect	12
Index	15

rdlocrand-package	<i>rdlocrand: Local Randomization Methods for RD Designs</i>
-------------------	--

Description

The regression discontinuity (RD) design is a popular quasi-experimental design for causal inference and policy evaluation. Under the local randomization approach, RD designs can be interpreted as randomized experiments inside a window around the cutoff. The rdlocrand package provides tools to analyze RD designs under local randomization: `rdrandinf` to perform hypothesis testing using randomization inference, `rdwinselect` to select a window around the cutoff in which randomization is likely to hold, `rdsensitivity` to assess sensitivity to different window lengths and null hypotheses, and `rdrbounds` to construct Rosenbaum bounds for sensitivity to unobserved confounders. For more details, and related R, Python, and Stata packages useful for analysis of RD designs, visit <https://rdpackages.github.io/>.

Author(s)

Matias D. Cattaneo, Princeton University. <matias.d.cattaneo@gmail.com>

Rocio Titiunik, Princeton University. <rocio.titiunik@gmail.com>

Gonzalo Vazquez-Bare, UC Santa Barbara. <gvazquezbare@gmail.com>

References

- Cattaneo, M.D., B. Frandsen and R. Titiunik. (2015). **Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate**. *Journal of Causal Inference* 3(1): 1-24.
- Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2016). **Inference in Regression Discontinuity Designs under Local Randomization**. *Stata Journal* 16(2): 331-367.
- Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2017). **Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality**. *Journal of Policy Analysis and Management* 36(3): 643-681.
- Rosenbaum, P. (2002). *Observational Studies*. Springer.

See Also

Useful links:

- <https://rdpackages.github.io/>
- <https://github.com/rdpackages/rdlocrand>
- Report bugs at <https://github.com/rdpackages/rdlocrand/issues>

rdrandinf

Randomization Inference for RD Designs under Local Randomization

Description

rdrandinf implements randomization inference and related methods for RD designs, using observations in a specified or data-driven selected window around the cutoff where local randomization is assumed to hold.

Usage

```
rdrandinf(  
  Y,  
  R,  
  cutoff = 0,  
  wl = NULL,  
  wr = NULL,  
  statistic = "diffmeans",  
  p = 0,  
  evall = NULL,  
  evalr = NULL,  
  kernel = "uniform",  
  fuzzy = NULL,  
  nulltau = 0,  
  d = NULL,  
  dscale = NULL,  
  ci,  
  interfci = NULL,  
  bernoulli = NULL,  
  reps = 1000,  
  seed = 666,  
  quietly = FALSE,  
  covariates,  
  obsmin = NULL,  
  wmin = NULL,  
  wobs = NULL,  
  wstep = NULL,  
  wasymmetric = FALSE,  
  wmasspoints = FALSE,
```

```

nwindows = 10,
dropmissing = FALSE,
rdwstat = "diffmeans",
approx = FALSE,
rdwreps = 1000,
level = 0.15,
plot = FALSE,
firststage = FALSE,
obsstep = NULL
)

```

Arguments

<code>Y</code>	a vector containing the values of the outcome variable.
<code>R</code>	a vector containing the values of the running variable.
<code>cutoff</code>	the RD cutoff (default is 0).
<code>wl</code>	the left limit of the window. The default takes the minimum of the running variable.
<code>wr</code>	the right limit of the window. The default takes the maximum of the running variable.
<code>statistic</code>	the randomization test statistic to be used. Allowed options are <code>diffmeans</code> (difference in means statistic), <code>ksmirnov</code> (Kolmogorov-Smirnov statistic), <code>ranksum</code> (Wilcoxon-Mann-Whitney standardized statistic), and <code>all</code> . Default option is <code>diffmeans</code> . The statistic <code>ttest</code> is equivalent to <code>diffmeans</code> and included for backward compatibility.
<code>p</code>	the order of the polynomial for the outcome adjustment model (default is 0).
<code>evall</code>	the point to the left of the cutoff at which the adjusted outcome is evaluated. Default is the cutoff value.
<code>evalr</code>	the point to the right of the cutoff at which the adjusted outcome is evaluated. Default is the cutoff value.
<code>kernel</code>	specifies the type of kernel to use as a weighting scheme. Allowed kernel types are <code>uniform</code> (uniform kernel), <code>triangular</code> (triangular kernel), and <code>epan</code> (Epanechnikov kernel). Default is <code>uniform</code> .
<code>fuzzy</code>	indicates that the RD design is fuzzy. <code>fuzzy</code> can be specified as a vector containing the values of the endogenous treatment variable, or as a list where the first element is the vector of endogenous treatment values and the second element is a string containing the statistic to be used. Allowed statistics are <code>ar</code> or <code>itt</code> (Anderson-Rubin/intention-to-treat statistic) and <code>tsls</code> (2SLS statistic). Default statistic is <code>ar</code> . The <code>tsls</code> statistic relies on a large-sample approximation.
<code>nulltau</code>	the value of the treatment effect under the null hypothesis (default is 0).
<code>d</code>	the effect size for asymptotic power calculation. Default is 0.5 times the standard deviation of the outcome variable for the control group.
<code>dscale</code>	the fraction of the standard deviation of the outcome variable for the control group used as the alternative hypothesis for asymptotic power calculation. Default is 0.5.

ci	calculates a confidence interval for the treatment effect by test inversion. ci can be specified as a scalar or a vector, where the first element indicates the value of alpha for the confidence interval (typically 0.05 or 0.01) and the remaining elements, if specified, indicate the grid of treatment effects to be evaluated. This option uses rdsensitivity to calculate the confidence interval. See the corresponding help file for details. Note: the default tlist can be narrow in some cases, which may truncate the confidence interval. We recommend manually setting a large enough tlist.
interfci	the level for Rosenbaum's confidence interval under arbitrary interference between units.
bernoulli	the probabilities of treatment for each unit when assignment mechanism is a Bernoulli trial. This option should be specified as a vector of length equal to the length of the outcome and running variables.
reps	the number of replications (default is 1000).
seed	the seed to be used for the randomization test.
quietly	suppresses the output table.
covariates	the covariates used by rdwinselect to choose the window when wl and wr are not specified. This should be a matrix of size n x k where n is the total sample size and k is the number of covariates.
obsmin	the minimum number of observations above and below the cutoff in the smallest window used by the companion command rdwinselect. Default is 10.
wmin	the smallest window to be used (if obsmin is not specified) by the companion command rdwinselect. Specifying both wmin and obsmin returns an error.
wobs	the number of observations to be added on each side of the cutoff at each step.
wstep	the increment in window length (if obsstep is not specified) by the companion command rdwinselect. Specifying both obsstep and wstep returns an error.
wasymmetric	allows for asymmetric windows around the cutoff when wobs is specified.
wmasspoints	specifies that the running variable is discrete and each masspoint should be used as a window.
nwindows	the number of windows to be used by the companion command rdwinselect. Default is 10.
dropmissing	drop rows with missing values in covariates when calculating windows.
rdwstat	the statistic to be used by the companion command rdwinselect (see the corresponding help file for options). Default option is diffmeans.
approx	forces the companion command rdwinselect to conduct the covariate balance tests using a large-sample approximation instead of finite-sample exact randomization inference methods.
rdwreps	the number of replications to be used by the companion command rdwinselect. Default is 1000.
level	the minimum accepted value of the p-value from the covariate balance tests to be used by the companion command rdwinselect. Default is .15.
plot	draws a scatter plot of the minimum p-value from the covariate balance test against window length implemented by the companion command rdwinselect.

firststage reports the results from the first step when using tsls.
 obsstep the minimum number of observations to be added on each side of the cutoff for the sequence of fixed-increment nested windows. Default is 2. This option is deprecated and only included for backward compatibility.

Value

A list containing:

sumstats matrix of full-sample and window-specific summary statistics.
 obs.stat observed statistic or statistics.
 p.value randomization p-value or p-values.
 asy.pvalue asymptotic p-value or p-values.
 window chosen window endpoints.
 ci confidence interval; included only when ci is specified.
 interf.ci confidence interval under interference; included only when interfci is specified.

Author(s)

Matias D. Cattaneo, Princeton University. <matias.d.cattaneo@gmail.com>

Rocio Titiunik, Princeton University. <rocio.titiunik@gmail.com>

Gonzalo Vazquez-Bare, UC Santa Barbara. <gvazquezbare@gmail.com>

References

Cattaneo, M.D., B. Frandsen and R. Titiunik. (2015). [Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate](#). *Journal of Causal Inference* 3(1): 1-24.

Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2016). [Inference in Regression Discontinuity Designs under Local Randomization](#). *Stata Journal* 16(2): 331-367.

Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2017). [Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality](#). *Journal of Policy Analysis and Management* 36(3): 643-681.

Examples

```
# Toy dataset
set.seed(123)
X <- array(rnorm(200),dim=c(100,2))
R <- X[,1] + X[,2] + rnorm(100)
Y <- 1 + R -.5*R^2 + .3*R^3 + (R>=0) + rnorm(100)
# Randomization inference in window (-.75,.75)
tmp <- rdrandinf(Y,R,wl=-.75,wr=.75,quietly=TRUE)
# Randomization inference in window (-.75,.75), all statistics
tmp <- rdrandinf(Y,R,wl=-.75,wr=.75,statistic='all',quietly=TRUE)
# Randomization inference with window selection
```

```
# Note: low number of replications to speed up process.
# The user should increase the number of replications.
tmp <- rdrandinf(Y,R,statistic='all',covariates=X,wmin=.5,wstep=.125,
                rdwreps=500,level=0,quietly=TRUE)
```

rdrbounds

*Rosenbaum bounds for RD designs under local randomization***Description**

rdrbounds calculates lower and upper bounds for the randomization p-value under different degrees of departure from a local randomized experiment, as suggested by Rosenbaum (2002).

Usage

```
rdrbounds(
  Y,
  R,
  cutoff = 0,
  wlist,
  gamma,
  expgamma,
  bound = "both",
  statistic = "ranksum",
  p = 0,
  evalat = "cutoff",
  kernel = "uniform",
  fuzzy = NULL,
  nulltau = 0,
  prob,
  fmpval = FALSE,
  reps = 1000,
  seed = 666
)
```

Arguments

Y	a vector containing the values of the outcome variable.
R	a vector containing the values of the running variable.
cutoff	the RD cutoff (default is 0).
wlist	the list of window lengths to be evaluated. By default the program constructs 10 windows around the cutoff, the first one including 10 treated and control observations and adding 5 observations to each group in subsequent windows.
gamma	the list of values of gamma to be evaluated.

expgamma	the list of values of exp(gamma) to be evaluated. Default is c(1.5, 2, 2.5, 3).
bound	specifies which bounds the command calculates. Options are upper for upper bound, lower for lower bound, and both for both upper and lower bounds. Default is both.
statistic	the randomization test statistic to be used. Allowed options are diffmeans (difference in means statistic), ksmirnov (Kolmogorov-Smirnov statistic), and ranksum (Wilcoxon-Mann-Whitney standardized statistic). Default option is ranksum. The statistic ttest is equivalent to diffmeans and included for backward compatibility.
p	the order of the polynomial for the outcome adjustment model. Default is 0.
evalat	specifies the point at which the adjusted variable is evaluated. Allowed options are cutoff and means. Default is cutoff.
kernel	specifies the type of kernel to use as a weighting scheme. Allowed kernel types are uniform (uniform kernel), triangular (triangular kernel), and epan (Epanechnikov kernel). Default is uniform.
fuzzy	indicates that the RD design is fuzzy. fuzzy should be specified as a vector containing the values of the endogenous treatment variable. This option uses an Anderson-Rubin/intention-to-treat statistic.
nulltau	the value of the treatment effect under the null hypothesis. Default is 0.
prob	the probabilities of treatment for each unit when the assignment mechanism is a Bernoulli trial. This option should be specified as a vector of length equal to the length of the outcome and running variables.
fmpval	reports the p-value under fixed margins randomization, in addition to the p-value under Bernoulli trials.
reps	the number of replications. Default is 1000.
seed	the seed to be used for the randomization tests.

Value

A list containing:

gamma	vector of gamma values.
expgamma	vector of exp(gamma) values.
wlist	window grid.
p.values	p-values for each window under gamma = 0. When fmpval = TRUE, this includes Bernoulli and fixed-margins p-values.
lower.bound	matrix of lower-bound p-values for each gamma-window pair; included when bound = "lower" or bound = "both".
upper.bound	matrix of upper-bound p-values for each gamma-window pair; included when bound = "upper" or bound = "both".

Author(s)

Matias D. Cattaneo, Princeton University. <matias.d.cattaneo@gmail.com>

Rocio Titiunik, Princeton University. <rocio.titiunik@gmail.com>

Gonzalo Vazquez-Bare, UC Santa Barbara. <gvazquezbare@gmail.com>

References

- Cattaneo, M.D., B. Frandsen and R. Titiunik. (2015). **Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate.** *Journal of Causal Inference* 3(1): 1-24.
- Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2016). **Inference in Regression Discontinuity Designs under Local Randomization.** *Stata Journal* 16(2): 331-367.
- Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2017). **Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality.** *Journal of Policy Analysis and Management* 36(3): 643-681.
- Rosenbaum, P. (2002). *Observational Studies*. Springer.

Examples

```
# Toy dataset
set.seed(123)
R <- runif(100,-1,1)
Y <- 1 + R - .5*R^2 + .3*R^3 + (R>=0) + rnorm(100)
# Rosenbaum bounds
# Note: low number of replications and windows to speed up process.
# The user should increase these values.
rdrbounds(Y,R,expgamma=c(1.5,2),wlist=c(.3),reps=100)
```

rdsensitivity

Sensitivity analysis for RD designs under local randomization

Description

rdsensitivity analyzes the sensitivity of randomization p-values and confidence intervals to different window lengths.

Usage

```
rdsensitivity(
  Y,
  R,
  cutoff = 0,
  wlist,
  wlist_left,
  tlist,
  statistic = "diffmeans",
  p = 0,
  evalat = "cutoff",
  kernel = "uniform",
  fuzzy = NULL,
```

```

ci = NULL,
ci_alpha = 0.05,
reps = 1000,
seed = 666,
nodraw = FALSE,
quietly = FALSE
)

```

Arguments

<code>Y</code>	a vector containing the values of the outcome variable.
<code>R</code>	a vector containing the values of the running variable.
<code>cutoff</code>	the RD cutoff (default is 0).
<code>wlist</code>	the list of windows to the right of the cutoff. By default the program constructs 10 windows around the cutoff with 5 observations each.
<code>wlist_left</code>	the list of windows to the left of the cutoff. If not specified, the windows are constructed symmetrically around the cutoff based on the values in <code>wlist</code> .
<code>tlist</code>	the list of treatment-effect values under the null to be evaluated. By default the program uses ten evenly spaced points within the asymptotic confidence interval for a constant treatment effect in the smallest window to be used.
<code>statistic</code>	the randomization test statistic to be used. Allowed options are <code>diffmeans</code> (difference in means statistic), <code>ksmirnov</code> (Kolmogorov-Smirnov statistic), and <code>ranksum</code> (Wilcoxon-Mann-Whitney standardized statistic). Default option is <code>diffmeans</code> . The statistic <code>ttest</code> is equivalent to <code>diffmeans</code> and included for backward compatibility.
<code>p</code>	the order of the polynomial for the outcome adjustment model. Default is 0.
<code>evalat</code>	specifies the point at which the adjusted variable is evaluated. Allowed options are <code>cutoff</code> and <code>means</code> . Default is <code>cutoff</code> .
<code>kernel</code>	specifies the type of kernel to use as a weighting scheme. Allowed kernel types are <code>uniform</code> (uniform kernel), <code>triangular</code> (triangular kernel), and <code>epan</code> (Epanechnikov kernel). Default is <code>uniform</code> .
<code>fuzzy</code>	indicates that the RD design is fuzzy. <code>fuzzy</code> should be specified as a vector containing the values of the endogenous treatment variable. This option uses an Anderson-Rubin/intention-to-treat statistic.
<code>ci</code>	returns the confidence interval corresponding to the indicated window length. <code>ci</code> must be a two-element vector containing the left and right limits of the window. Default alpha is .05 (95% level CI).
<code>ci_alpha</code>	specifies the value of alpha for the confidence interval. Default alpha is .05 (95% level CI).
<code>reps</code>	the number of replications. Default is 1000.
<code>seed</code>	the seed to be used for the randomization tests.
<code>nodraw</code>	suppresses contour plot.
<code>quietly</code>	suppresses the output table.

`rdwinselect`*Window selection for RD designs under local randomization*

Description

`rdwinselect` implements a window-selection procedure based on balance tests for RD designs under local randomization. Specifically, it constructs a sequence of nested windows around the RD cutoff and reports binomial tests for the running variable and covariate balance tests for the covariates (if specified). The recommended window is the largest window around the cutoff such that the minimum p-value from the balance tests is larger than a prespecified level for all nested (smaller) windows. By default, the p-values are calculated using randomization inference methods.

Usage

```
rdwinselect(  
  R,  
  X,  
  cutoff = 0,  
  obsmin = NULL,  
  wmin = NULL,  
  wobs = NULL,  
  wstep = NULL,  
  wasymmetric = FALSE,  
  wmasspoints = FALSE,  
  dropmissing = FALSE,  
  nwindows = 10,  
  statistic = "diffmeans",  
  p = 0,  
  evalat = "cutoff",  
  kernel = "uniform",  
  approx = FALSE,  
  level = 0.15,  
  reps = 1000,  
  seed = 666,  
  plot = FALSE,  
  quietly = FALSE,  
  obsstep = NULL  
)
```

Arguments

R a vector containing the values of the running variable.

X the matrix of covariates to be used in the balance tests. The matrix is optional, but the recommended window is only provided when at least one covariate is specified. This should be a matrix of size $n \times k$ where n is the total sample size and k is the number of covariates.

cutoff	the RD cutoff (default is 0).
obsmin	the minimum number of observations above and below the cutoff in the smallest window. Default is 10.
wmin	the smallest window to be used.
wobs	the number of observations to be added on each side of the cutoff at each step. Default is 5.
wstep	the increment in window length.
wasymmetric	allows for asymmetric windows around the cutoff when wobs is specified.
wmasspoints	specifies that the running variable is discrete and each masspoint should be used as a window.
dropmissing	drop rows with missing values in covariates when calculating windows.
nwindows	the number of windows to be used. Default is 10.
statistic	the statistic to be used in the balance tests. Allowed options are <code>diffmeans</code> (difference in means statistic), <code>ksmirnov</code> (Kolmogorov-Smirnov statistic), <code>ranksum</code> (Wilcoxon-Mann-Whitney standardized statistic) and <code>hotelling</code> (Hotelling's T-squared statistic). Default option is <code>diffmeans</code> . The statistic <code>ttest</code> is equivalent to <code>diffmeans</code> and included for backward compatibility.
p	the order of the polynomial for the outcome adjustment model (for covariates). Default is 0.
evalat	specifies the point at which the adjusted variable is evaluated. Allowed options are <code>cutoff</code> and <code>means</code> . Default is <code>cutoff</code> .
kernel	specifies the type of kernel to use as a weighting scheme. Allowed kernel types are <code>uniform</code> (uniform kernel), <code>triangular</code> (triangular kernel), and <code>epan</code> (Epanechnikov kernel). Default is <code>uniform</code> .
approx	forces the command to conduct the covariate balance tests using a large-sample approximation instead of finite-sample exact randomization inference methods.
level	the minimum accepted value of the p-value from the covariate balance tests. Default is .15.
reps	the number of replications. Default is 1000.
seed	the seed to be used for the randomization tests.
plot	draws a scatter plot of the minimum p-value from the covariate balance test against window length.
quietly	suppresses output.
obsstep	the minimum number of observations to be added on each side of the cutoff for the sequence of fixed-increment nested windows. This option is deprecated and only included for backward compatibility.

Value

A list containing:

w_left	left endpoint of the recommended window.
w_right	right endpoint of the recommended window.

wlist_left	left endpoints of the candidate windows.
wlist_right	right endpoints of the candidate windows.
results	matrix containing the minimum covariate-balance p-value, selected covariate index, binomial-test p-value, sample sizes below and above the cutoff, and window endpoints for each candidate window.
summary	matrix of sample-size summaries by side of the cutoff.

Author(s)

Matias D. Cattaneo, Princeton University. <matias.d.cattaneo@gmail.com>

Rocio Titiunik, Princeton University. <rocio.titiunik@gmail.com>

Gonzalo Vazquez-Bare, UC Santa Barbara. <gvazquezbare@gmail.com>

References

Cattaneo, M.D., B. Frandsen and R. Titiunik. (2015). **Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate.** *Journal of Causal Inference* 3(1): 1-24.

Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2016). **Inference in Regression Discontinuity Designs under Local Randomization.** *Stata Journal* 16(2): 331-367.

Cattaneo, M.D., R. Titiunik and G. Vazquez-Bare. (2017). **Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality.** *Journal of Policy Analysis and Management* 36(3): 643-681.

Examples

```
# Toy dataset
set.seed(123)
X <- array(rnorm(200),dim=c(100,2))
R <- X[,1] + X[,2] + rnorm(100)
# Window selection adding 5 observations at each step
# Note: low number of replications to speed up process.
tmp <- rdwinselect(R,X,obsmin=10,wobs=5,nwindows=5, reps=500,quietly=TRUE)
# Window selection setting initial window and step
# The user should increase the number of replications.
tmp <- rdwinselect(R,X,wmin=.5,wstep=.125, reps=500,quietly=TRUE)
# Window selection with approximate (large sample) inference and p-value plot
tmp <- rdwinselect(R,X,wmin=.5,wstep=.125, approx=TRUE, nwindows=20,quietly=TRUE,plot=TRUE)
```

Index

`rdlocrand` (`rdlocrand-package`), [2](#)
`rdlocrand-package`, [2](#)
`rdlocrand_package` (`rdlocrand-package`), [2](#)
`rdrandinf`, [2](#), [3](#)
`rdrbounds`, [2](#), [7](#)
`rdsensitivity`, [2](#), [9](#)
`rdwinselect`, [2](#), [12](#)